Efficient HTTP based I/O on very large datasets for high performance computing with the libdavix library

Adrien Devresse and Fabrizio Furano

CERN, European Organization for Nuclear Research ,
Geneva, Switzerland
adrien.devresse@cern.ch,furano@cern.ch
http://cern.ch

arXiv:1410.4168v1 [cs.PF] 15 Oct 2014

Abstract. Remote data access for data analysis in high performance computing is commonly done with specialized data access protocols and storage systems. These protocols are highly optimized for high throughput on very large datasets, multi-streams, high availability, low latency and efficient parallel I/O. The purpose of this paper is to describe how we have adapted a generic protocol, the Hyper Text Transport Protocol (HTTP) to make it a competitive alternative for high performance I/O and data analysis applications in a global computing grid: the Worldwide LHC Computing Grid. In this work, we first analyze the design differences between the HTTP protocol and the most common high performance I/O protocols, pointing out the main performance weaknesses of HTTP. Then, we describe in detail how we solved these issues. Our solutions have been implemented in a toolkit called davix, available through several recent Linux distributions.
Finally, we describe the results of our benchmarks where we compare the performance of davix against a HPC specific protocol for a data analysis use case.

Keywords: High Performance Computing, Big Data, HTTP Protocol, Data Access, Performance Optimization, I/O

1 Introduction

The Hyper Text Transport Protocol (HTTP)[9] is today undoubtedly one of the most prevalent protocols on the internet.

Initially created by Tim Berners-Lee for the World Wide Web at CERN in 1990[8], HTTP is today much more than a simple protocol dedicated to HTML web page transport. The extensions of HTTP like WebDav[21] or CalDav[15], the HTTP based protocols like UPnP or SOAP[10] and the RESTful[17] architecture for Web Services have transformed HTTP into an universal versatile application layer protocol for the internet. The recent emergence of cloud computing[13] and the popularization of big data storage[31] based on RESTful Web services have definitively proved the universalism of HTTP.

Today, HTTP is the foundation for interactions with commercial cloud storage providers like Amazon Simple Storage Service[27] or with Open Source Cloud Storage[5] systems like OpenStack Swift [29] using REST API like S3[26] or CDMI[36]. HTTP is fully accepted as data transfer and data manipulation protocol in NoSQL databases and distributed storage systems in the Web World. The association of REST APIs with the HTTP protocol usage offers a simple, standard, extensible, portable alternative to the legacy data access and file manipulation protocols or to the proprietary protocol of most distributed file systems.

However, the popularity of HTTP was still not penetrating the High Performance Computing world. HPC[1] data access have highly specific and strict requirements: very large data manipulation, low-latency, high-throughput, high-availability, highly parallel I/O and high-scalability. For these reasons, those use cases traditionally rely on highly specific systems and protocols adapted to such constrains. The IBM GPFS[35] protocol, the Lustre parallel distributed file system protocol, the Hadoop HDFS[42] streaming protocol, the gridFTP protocol[11] or the XRootD protocol[31] are widely used in super computing and grid computing environments.

The focus of this work is to be able to make the HPC world benefit of all the momentum coming from the HTTP Ecosystem, like the RESTful and Cloud Storage services, by creating a high performance I/O layer based on the HTTP protocol, compatible with standard services and competitive with the HPC I/O specific protocols.
To achieve this, we have created *libdavix*[38][1], an I/O layer implementation optimized for data analysis and HPC I/O in distributed environments.

2 Background and Related work

2.1 HTTP as a data management and data transfer protocol

The stateless nature of HTTP, associated with the atomic nature of its primitives, provides a simple, reliable and powerful consistency model in distributed environment. The success and the major diffusion of the RESTful architecture introduced by Roy Fielding[17] for the Web World illustrate perfectly this [34].

The HTTP PUT method provides an object level idempotent write operation that can be used for an atomic resource creation or a resource content update. The basic HTTP GET method gives a safe, cacheable, atomic and idempotent Object level read operation and can be used to access and retrieve safely a remote resource. These two methods, associated with the HTTP DELETE method, are enough to satisfy the four basic functions Create, Retrieve, Update, Delete (CRUD)[32] of any basic persistent storage system [6].

The properties of the HTTP protocol make it suitable for data transfer in a distributed environment and easily justify the emergence of persistency and data storage solutions using RESTful interfaces. It is the case for instance of the

[1] High Performance Computing

NoSQL database couchdb[3], of the NoSQL key-store Ryak[2], of the distributed file system HDFS with httpFS[42], of the Amazon Simple Storage Service (Amazon S3)[27] or of any similar RESTful Object Storage service.
Again, the universalism of HTTP and the quality of its ecosystem associated with its simple and flexible design, makes it a first quality choice for a generic data transfer protocol today.

2.2 Efficient parallel I/O operations

Intensive data analysis applications requires high degree of I/O parallelism, robustness over large transfers and low I/O latency. A high energy physics application typically processes in parallel a very large number of events from different files located in large distributed data stores, triggering a significant number of parallel I/O operations.

For such use cases in, the grid computing models in the HEP community use a mix of I/O frameworks for HPC[2] data access, the XRootD framework[16], the GridFTP protocol[11] with the Globus toolkit[18], HDFS of Hadoop[42] and IBM GPFS[35].
All these frameworks are highly optimized for parallel access, high throughput and efficient I/O scheduling of multiple requests. The XRootD framework implements its own I/O scheduler, it supports parallel asynchronous data access on top of its own I/O multiplexing.
The GridFTPv2 protocol has separated control and data channels and supports multiple data streams from different data sources on top of TCP or UDP.
The HDFS architecture is specially designed for large file storage, high throughput, hot file replications and data streams from multiple DataNodes.

To explore the possibilities for a solution based on HTTP, we defined a set of criteria to meet.

- Efficiency for large data transport and parallel I/O.
- Compatibility with existing network infrastructure and services.
- Low I/O latency: avoid useless handshakes, useless reconnections and redirections.

The original design of HTTP did not match very well these points.
The HTTP 1.0 standard recommends the usage of one TCP connection per request to the server. This approach has been already proven inefficient due to the TCP slow start mechanism [37]. Executing a HEP[3] data analysis work-flow with a very large number of parallel small sized requests in such conditions would lead to a major performance penalty.

To mitigate the effect of this behaviour, HTTP 1.1 introduced the persistent connection support with the *KeepAlive* mechanism and the support of request pipelining over the same connection [34].

[2] High performance computing
[3] High energy physics

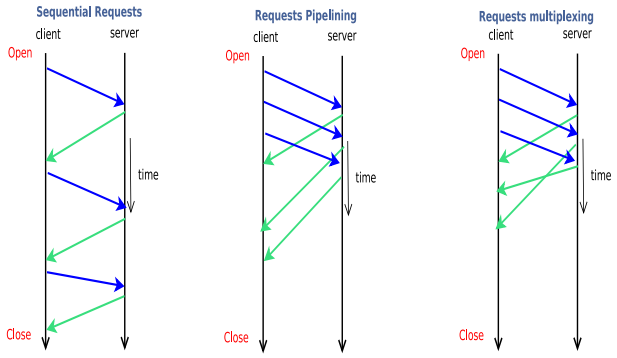

Fig. 1. HTTP Request pipelining and Request Multiplexing

However, HTTP pipelining suffers of several problems. Contrary to a protocol supporting modern multiplexing, the HTTP pipelining suffers of the HOL[4] problem. The HTTP standard specifies that the treatment of a group of pipelined HTTP requests has to be processed in order. With such a requirement, any request pipelined suffering of a delay will cause a delay for all the following requests [28]. This is a unacceptable performance penalty in case of parallel I/O requests with different sizes.
The HTTP pipelining also suffers from other problems. It suffers of side effects with the TCP's nagle algorithm[22]. It often suffers of performance degradation due to aggressive pipeline interruptions with some servers implementations and due to the fact that the pipelining is not respected by most of the proxy servers. For these reasons, most of the current web browsers (Chrome, Firefox and Internet explorer) web browser disable or do not support the HTTP pipelining mechanism.

To resolve these problems inherent to HTTP 1.1, several proposal have been made:

- the **SPDY protocol** is "an application-layer protocol for transporting content over the web, designed specifically for minimal latency". SPDY acts as a session layer between HTTP and TCP. It supports multiplexing, prioritization and header compression[7]. SPDY is currently the most mature implementation of multi-plexing for HTTP and supported by a majority of modern web browsers.
- **HTTP over SCTP** proposes to use the SCTP multi-homing and multiplexing features with HTTP[33]. SCTP is a acronym for Stream Control Transmission Protocol, it provides a message-based alternative to TCP.
- **WebMux and HTTP-NG** proposes the usage of the MUX protocol as session protocol to provide multi-plexing for HTTP[30][20].

[4] Head of Line

None of these approaches are considered acceptable for a High performance I/O usage.

The SPDY protocol explicitly enforces the usage of SSL/TLS for protocol negotiation purpose. TLS introduces a negative performance impact for big data transfers [14] and introduces a handshake latency that can not be mandatory in High performance computing.

The HTTP over SCTP proposal implies naturally to replace the TCP protocol by SCTP. Like any other protocol aiming to replace the level 4 of the OSI model, the SCTP protocol triggers several concerns about the compatibility with the existing network architecture, about the NAT-traversals capability and about the support in old operating systems.

The MUX, renamed WebMUX protocol, defined in 1998, focuses on an object oriented approach with in mind the support for technology like RMI[5], DCOM[6] or CORBA[7] which is not our use cases. Moreover, it has currently not been implemented in any major HTTP server.

To satisfy our use cases, we adopted and implemented a different approach into *libdavix* [38][1]. We use a hybrid solution based on a dynamic connection pool with a thread-safe query dispatch system and a session recycling mechanism (See Figure-2.2).
Associated with the pool, we enforce an aggressive usage of the HTTP KeepAlive feature, *libdavix* to maximize the re-utilization of the TCP connections and to minimize the effect of the TCP slow start.

This method gives several benefits compared to previously quoted solutions. First, it is fully compatibility with the standard HTTP 1.1 and with existing services and infrastructure.
Second, we supports efficient parallel request execution for repetitive I/O operations without suffering of the problems that are specific to the classical HTTP pipelining. nor necessitating a protocol modification to support multi-plexing.

This dispatch approach is particularly adapted to a HPC data-analysis workflow: a repetitive concurrent access to a limited set of hosts exploits at the best the session recycling and maximizes the lifetime of the TCP connections. However, contrary to a pure multi-plexing solution that aims to the usage of one TCP connection per host, our approach uses a connection pool whose size is proportional to the level of concurrency. Consequently, an important degree of concurrency can result in a more important server load compared to a multiplexed solution like spdy due to the number of connections allocated per client.

However, this is not a big issue for our HPC use case, the support of "vectored queries" of *libdavix* explained in section 2.3 reducing significantly the number of simultaneous concurrent requests.

[5] Remote Method Invocation
[6] Distributed Component Object Model
[7] Common Object Request Broker Architecture

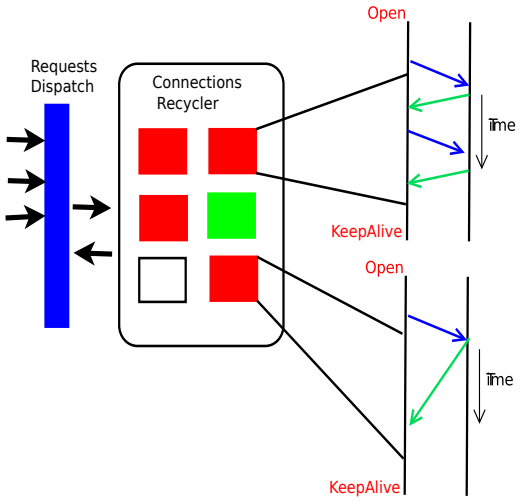

Fig. 2. Dynamic connection pool with a thread-safe query dispatch

2.3 Scalable random-access I/O for partial file access with HTTP

Very large data sets in distributed environments are generally partitioned into separated data subset objects distributed in several storage nodes.
In High Energy Physics, a data set generally contains a important number of particle events decomposed in ROOT[4] *TTrees* of events and stored in different compressed files.

This approach allows an easy data distribution and simplify data replication. It facilitates the partitioning the data in a distributed environment like in the the World Wide LHC Computing Grid (WLCG)[23]. At the same time, a data analysis with this data model is I/O intensive and generates a very large number of individual data accesses operation to the storage. In order to extract a set of specific events spread in different remote data sets, a HEP Application needs to read a large number of small segments of data in different remote objects.
To reduce the number of I/O requests, and, hence, the impact of latency with such patterns, high performance computing I/O protocols implement Data Sieving algorithms, two phase I/O algorithms[39] or sliding window buffering algorithms.
To the best of our knowledge, no nowadays HTTP I/O toolkit before davix implemented similar I/O optimization strategies.

We implemented in *libdavix* a support for vectored packed operations with random position based on the multi-range feature of the HTTP 1.1 protocol (Figure 2.3).
This feature allows to gather and pack a large number of fragmented random

I/O requests directly in the ROOT[4] I/O framework via the TTreeCache[41] in a large vectored query. Subsequently, this query is processed by *libdavix* as one atomic remote I/O query. This approach reduces drastically the number of remote network I/O operations and offers the advantage to reduce the necessity of parallel I/O operations and thus virtually eliminates the need for I/O multiplexing.

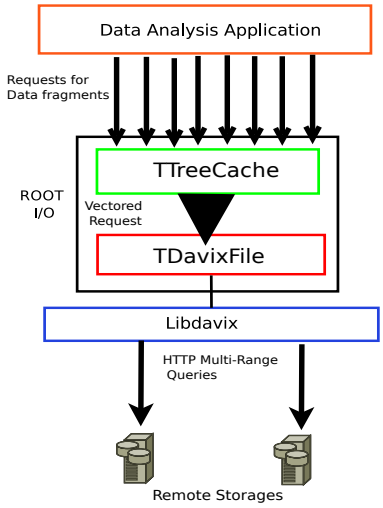

Fig. 3. Vectored I/O requests support in the ROOT framework associated with LibDavix

2.4 Multi-stream and multiple replicas I/O operations

In grid computing computing, the unavailability of an input data required by a job is often the main cause of failure. Such situation leads to a redistribution of the data and a rescheduling of the job with a substantial performance impact on the execution time.

In an attempt to solve this issue, the XRootD protocol [16] supports a federation mechanism to offer a more resilient access to a distributed resources. XRootD data servers can be federated hierarchically into a global virtual namespace. In case of unavailability of a resources in the closest data repository, the XRootD federation mechanism will locate a second available replica of this resource and redirect the client there.

HTTP alone does not support this feature. A classical HTTP access to an unavailable resource or a offline server will result in a I/O error. To improve the resiliency of the data layer, davix implements natively a support for the

Metalink[25] standard file format coupled to a fail-over and filtering mechanism for the offline replicas of a resource.

A Metalink file is a standardized XML[12] file containing several elements of meta-data information about an online resource: name, size, checksum, signature and location of the replicas of the resource. A Metalink file is a resource description and a set of ordered pointers to this resource.

Davix can use the Metalink information to apply two strategies:

- The **"fail-over" strategy** (default). In the case a resource is not available, davix try seamlessly to obtain the Metalink associated with this resource. Then *libdavix* will try to access one per one the available replicas of this resource until being able to access the requested data on one of the available replica of the resource. This approach improves drastically the resiliency of the data access layer and has the advantage to be without compromise or impact on the performances.
- The **"Multi-stream" strategy**. In this case, *libdavix* will first try to obtain the Metalink of the resource and then proceed to a multi-source parallel download of each referenced chunk of data from a different replica. This approach has the advantage to maximize the network bandwidth usage on the client side and to offer the same resiliency improvement than the fail-over strategy. However, it has for main drawback to overload considerably the servers.

The combined usage of *libdavix* for data analysis with a Replica catalogue or federation system supporting able to provide Metalink files like the DynaFed system (Dynamic Storage Federation)[19] enforces the global resilience of the I/O layer of any HPC application in a transparent manner. It provides the guarantee that a read operation on a resource will succeed as long as one replica of this resource is remotely accessible and referenced by the corresponding Metalink.

3 Performance analysis

For our performance analysis, we executed a High energy analysis job based on ROOT framework[4] reading a fraction or the totality of around 12000 particles events from a 700 MBytes root file. This tests has been executed using both the XRootD framework and our libdavix solution as I/O layer.

Each test execution has been executed on WLCG through the Hammerloud Grid[40] performance testing framework.
The execution of the test is always done on a standard Worder Node configuration of WLCG.
For both XRootD and davix, each test is run against the same server instance with the following configuration: Disk Pool Manager(DPM) Storage system, 4 Core Intel Xeon CPU, 12 GBytes of RAM, Scientific Linux CERN 6, 1 GB/s network link.

Two test executions are separated by 20 minutes.
Our statistics are based on an average of 576 tests executions over a period of 12 days.

We compare here the global execution time of the analysis job over different network configurations:

- **LAN**: accessing the file over a gigabit Ethernet with low latency (latency <5 ms)
- **PAN-European network**: accessing the file over GEANT[24] between Switzerland and UK (latency <50 ms)
- **WAN**: accessing the file over Internet between Switzerland and USA (latency <300 ms)

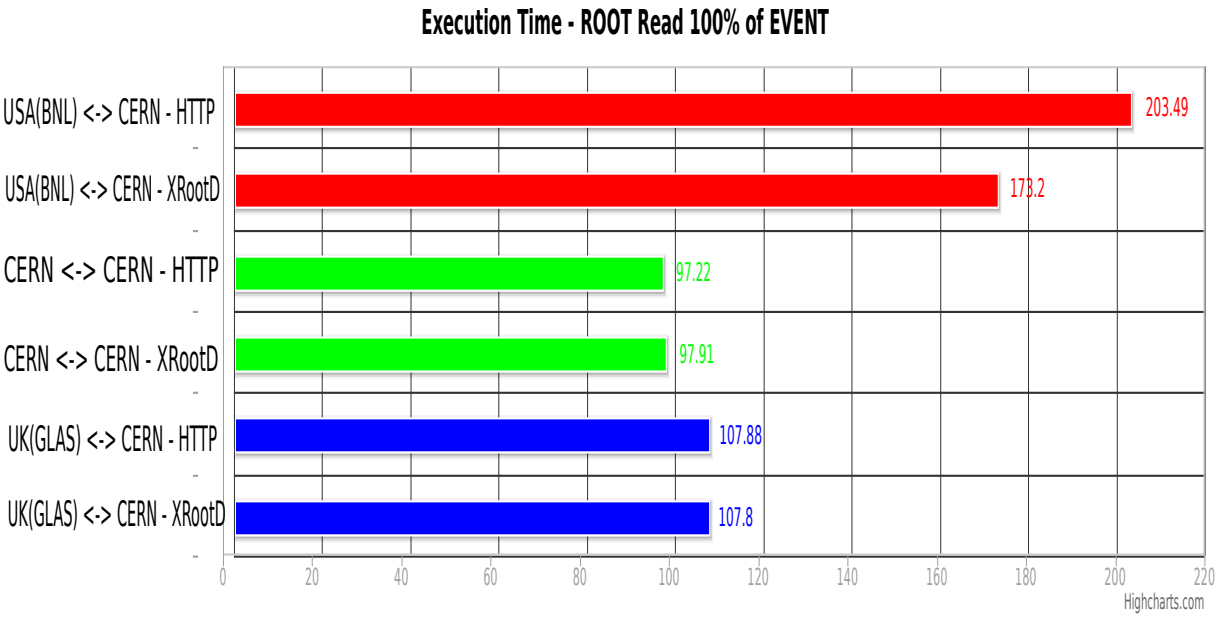

Fig. 4. Execution time of a ROOT data analysis job (less is better).

In case of "CERN-CERN" data transfer, the (Figure-3) shows that *libdavix* is respectively 0.7 % faster than XRootD in case of local access with high speed link and low latency. This shows that HTTP with *libdavix* can compete with a HPC specific I/O protocol on local area network and offers similar performances in term of data access time and data transfer rate.

In case of "UK(GLAS)-CERN", XRootd and *libdavix* offers sensibly the same performance.

In case of "USA(BNL)-CERN" data transfer, our tests shows that XRootD is in average 17.5% faster than *libdavix* on Wide Area network links with high latency. This difference of performance comes mainly from the sliding windows buffering algorithm of XRootD which allows to minimize the number of network round trips executed. Network round trips are naturally extremely costly on high latency networks.

In a classical High Energy Physics grid computing model, a job is always sent close to the data that it will process. Data access are in this case made over a LAN with high speed and low latency.
In such model, these results are particularly encouraging for *libdavix* and HTTP I/O.

4 Conclusion

Our results shows that for a HPC I/O use case, our solution, *libdavix* can provide similar performance over low latency link to a HPC specific protocol like XRootD.
The usage of the HTTP multi-range feature allows to reduce drastically the number of parallel I/O operations and allows HTTP to compete with the aggressive caching strategy of the HPC specific protocols in case remote partial I/O operations on large data sets.
The lack of multi-plexing support in HTTP can be compensated by a session recycling system for HEP uses cases and allows to be retro-compatible with the existing network and service infrastructure.
Finally, the association of the Metalink with HTTP gives new possibilities for transparent error recovery in HPC data access and offers an interesting alternative to classical hierarchical data federations.

We have demonstrated in this paper with *libdavix* that an optimized I/O layer based on the HTTP protocol can be considered as a serious and performant alternative to the common HPC specific I/O protocols.

5 Acknowledgments

We thank Oliver Keeble, Martin Hellmich, Ivan Calvet and Alejandro Alvarez Ayllon for their contributions, commitment and testing dedicated the davix toolkit. Thank you also to Tigran Mkrtchyan for his contributions to TDavixFile. We also thank the ROOT development team for their support and help during the integration of davix to the ROOT analysis framework. Thank you to Olivier Perrin for his support and advises in this work.

References

1. Davix project, doi:10.5281/zenodo.10761.
2. Riak: An open source scalable data store, 5 Juin 2014.
3. J. C. Anderson, J. Lehnardt, and N. Slater. *CouchDB: the definitive guide.* " O'Reilly Media, Inc.", 2010.
4. I. Antcheva, M. Ballintijn, B. Bellenot, M. Biskup, R. Brun, N. Buncic, P. Canal, D. Casadei, O. Couet, V. Fine, L. Franco, G. Ganis, A. Gheata, D. G. Maline, M. Goto, J. Iwaszkiewicz, A. Kreshuk, D. M. Segura, R. Maunder, L. Moneta, and A. Naumann. Root - a c++ framework for petabyte data storage, statistical analysis and visualization. *Computer Physics Communications*, 182(6):1384–1385, 2011.
5. S. A. Baset. Open source cloud technologies. In *Proceedings of the Third ACM Symposium on Cloud Computing*, SoCC '12, pages 28:1–28:2, New York, NY, USA, 2012. ACM.
6. R. Battle and E. Benson. Bridging the semantic web and web 2.0 with representational state transfer (rest). *Web Semant.*, 6(1):61–69, Feb. 2008.
7. M. Belshe and R. Peon. Spdy protocol. 2012.
8. T. Berners-Lee and R. Cailliau. Worldwideweb: Proposal for a hypertext project. *Retrieved on February*, 26:2008, 1990.
9. T. Berners-Lee, R. Fielding, and H. Frystyk. Hypertext transfer protocol–http/1.0, 1996.
10. D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. F. Nielsen, S. Thatte, and D. Winer. Simple object access protocol (soap) 1.1, 2000.
11. J. Bresnahan, M. Link, G. Khanna, Z. Imani, R. Kettimuthu, and I. Foster. Globus gridftp: What's new in 2007. In *Proceedings of the First International Conference on Networks for Grid Applications*, GridNets '07, pages 19:1–19:5, ICST, Brussels, Belgium, Belgium, 2007. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
12. A. Bryan, T. Tsujikawa, and N. McNab. P. poeml," the metalink download description format. Technical report, RFC 5854, June, 2010.
13. R. Buyya, J. Broberg, and A. Goscinski. *Cloud Computing: Principles and Paradigms.* Wiley Series on Parallel and Distributed Computing. Wiley, 2010.
14. C. Coarfa, P. Druschel, and D. S. Wallach. Performance analysis of tls web servers. *ACM Trans. Comput. Syst.*, 24(1):39–69, Feb. 2006.
15. C. Daboo, B. Desruisseaux, and L. Dusseault. Calendaring extensions to webdav (caldav). Technical report, RFC 4791, March, 2007.
16. A. Dorigo, P. Elmer, F. Furano, and A. Hanushevsky. Xrootd-a highly scalable architecture for data access. *WSEAS Transactions on Computers*, 1(4.3), 2005.
17. R. T. Fielding. *Architectural Styles and the Design of Network-based Software Architectures.* PhD thesis, 2000. AAI9980887.
18. I. Foster and C. Kesselman. The globus toolkit. *The grid: blueprint for a new computing infrastructure*, pages 259–278, 1999.
19. F. Furano, R. B. da Rocha, A. Devresse, O. Keeble, A. Á. Ayllón, and P. Fuhrmann. Dynamic federations: storage aggregation using open tools and protocols. In *Journal of Physics: Conference Series*, volume 396, page 032042. IOP Publishing, 2012.
20. J. Gettys and H. F. Nielsen. Smux protocol specification. *Work In Progress (W3C Working Draft WD-mux-19980710)*, 1998.
21. Y. Goland, E. Whitehead, A. Faizi, S. Carter, and D. Jensen. Http extensions for distributed authoring–webdav, 1999.

22. J. Heidemann. Performance interactions between p-http and tcp implementations. *SIGCOMM Comput. Commun. Rev.*, 27(2):65–73, Apr. 1997.
23. http://wlcg.web.cern.ch/. Worldwide lhc computing grid website.
24. http://www.geant.net. Geant project website.
25. http://www.metalinker.org/. Metalink project website.
26. A. Inc. Amazon s3 rest, 2014.
27. A. Inc. Amazon ss, 2014.
28. G. Inc. Spdy frequently asked questions, doesn't http pipelining already solve the latency problem ?, 2014.
29. K. Jackson. *OpenStack Cloud Computing Cookbook*. Packt Publishing Ltd, 2012.
30. B. Janssen. Http-ng overview problem statement, requirements, and solution outline.
31. J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. H. Byers. Big data: The next frontier for innovation, competition, and productivity. 2011.
32. J. Martin. *Managing the Data Base Environment*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1st edition, 1983.
33. P. Natarajan, J. R. Iyengar, P. D. Amer, and R. Stewart. Sctp: An innovative transport layer protocol for the web. In *Proceedings of the 15th International Conference on World Wide Web*, WWW '06, pages 615–624, New York, NY, USA, 2006. ACM.
34. H. RFC. Hypertext transport protocol rfc, 1999.
35. F. Schmuck and R. Haskin. Gpfs: A shared-disk file system for large computing clusters. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*, FAST '02, Berkeley, CA, USA, 2002. USENIX Association.
36. snia. Cdmi specification, 2014.
37. S. E. Spero. Analysis of http performance problems, July 1994.
38. D. team. Davix project website, http://dmc.web.cern.ch/projects/davix/home.
39. R. Thakur, W. Gropp, and E. Lusk. Data sieving and collective i/o in romio. In *Frontiers of Massively Parallel Computation, 1999. Frontiers '99. The Seventh Symposium on the*, pages 182–189, Feb 1999.
40. D. C. van der Ster, J. Elmsheuser, M. U. Garcia, and M. Paladin. Hammercloud: A stress testing system for distributed analysis. In *Journal of Physics: Conference Series*, volume 331, page 072036. IOP Publishing, 2011.
41. I. Vukotic, A. Collaboration, et al. Optimization and performance measurements of root-based data formats in the atlas experiment. In *Journal of Physics: Conference Series*, volume 331, page 032032. IOP Publishing, 2011.
42. T. White. *Hadoop: The Definitive Guide: The Definitive Guide*. O'Reilly Media, 2009.